

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
3404 E. Harmony Road
Mail Stop 35
Fort Collins, Colorado 80528

PATENT APPLICATION

ATTORNEY DOCKET NO. 200313629-1

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor(s): Simon C. Steely Jr. et al.

Confirmation No.: 9866

Application No.: 10/760,659

Examiner: Mardochee Chery

Filing Date: January 20, 2004

Group Art Unit: 2186

Title: System and Method for Non-Migratory Requests in a Cache Coherency Protocol

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF

Transmitted herewith is the Appeal Brief in this application with respect to the Notice of Appeal filed on February 10, 2010.

- ☒ The fee for filing this Appeal Brief is \$540.00 (37 CFR 41.20). minus \$500 for Appeal Brief
☐ No Additional Fee Required. previously filed on 9-12-07.

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

- ☐ (a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d)) for the total number of months checked below:

☐ 1st Month
\$130

☐ 2nd Month
\$490

☐ 3rd Month
\$1110

☐ 4th Month
\$1730

- ☐ The extension fee has already been filed in this application.

- ☒ (b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account 08-2025 the sum of \$ 40.00. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees.

Respectfully submitted,

Simon C. Steely Jr. et al.

By: /Steven L. Nichols/

Steven L. Nichols

Attorney/Agent for Applicant(s)

Reg No. : 40,326

Date : February 10, 2010

Telephone : 801-572-8066

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In the Patent Application of

Simon C. Steely, Jr. et al.

Application No. 10/760,659

Filed: January 20, 2004

For: System and Method for Non-Migratory
Requests in a Cache Coherency Protocol

Group Art Unit: 2186

Examiner: Chery, Mardochee

Confirmation No.: 9866

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Following Appellants' filing of an Appeal Brief on 8 February 2008, prosecution of this application was reopened in response to Appellant's submission of an Information Disclosure Statement together with a Request for Continued Examination on 13 August 2009. The Examiner subsequently issued a final Office Action dated 10 December 2009 (the "Office Action" or the "Action"). Having reviewed the final Office Action of December 10, 2009, Appellants hereby request re-instatement of the appeal in this application and file the present updated Appeal Brief in support of the re-instated appeal. Each of the topics required by Rule 41.37 is presented herewith and is labeled appropriately.

I. Real Party in Interest

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 11445 Compaq Center Drive W., Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

II. Related Appeals and Interferences

An appeal is pending for U.S. Patent Application Serial No. 10/760,652 which may be considered related to the present appeal.

III. Status of Claims

Claims 2-7 and 13-17 have been identified as reciting allowable subject matter, but stand objected to as being dependent upon a rejected base claim.

Claims 1, 8-12, and 18-32 are pending in the application and stand finally rejected.

Accordingly, Appellant appeals from the final rejection of claims 1, 8-12, and 18-32, which claims are presented in the Appendix.

IV. Status of Amendments

No amendments have been filed subsequent to the Office Action of December 10, 2009, from which Appellant takes this appeal.

V. Summary of Claimed Subject Matter

The claimed subject matter is directed to systems in which a first node (12, Fig. 1) includes data having an associated state that is a modified state. (*E.g.*, Appellant's Specification, paras. 0020-23). A second node (14, Fig. 1) is operative to provide a non-migratory source broadcast request for the data. (*E.g.*, Appellant's Specification, para. 0044). The first node (12, Fig. 1) is operative in response to the non-migratory source broadcast request to provide the data to the second node (14, Fig. 1). (*E.g.*, Appellant's Specification, para. 0044). The first node (12, Fig. 1) is further operative to transition the associated state of the data at the first node (12, Fig. 1) from the modified state to an owner state without updating memory. (*E.g.*, Appellant's Specification, para. 0044). The second node (14, Fig. 1) is operative to receive the data from the first node (12, Fig. 1) and assign a shared state to an associated state of the data at the second node (14, Fig. 1). (*E.g.*, Appellant's Specification, paras. 0020-44).

Turning now to Appellant's specific claims,

Claim 1 recites:

A system (10, Fig. 1) comprising:

a first node (12, Fig. 1) including data having an associated state, the associated state of the data at the first node (12, Fig. 1) being a modified state (*e.g.*, Appellant's Specification, paras. 0020-23); and

a second node (14, Fig. 1) operative to provide a non-migratory source broadcast request for the data (*e.g.*, Appellant's Specification, para. 0044), the first node (12, Fig. 1) being operative in response to the non-migratory source broadcast request to provide the data

to the second node (14, Fig. 1) (*e.g.*, Appellant's Specification, para. 0044) and to transition the associated state of the data at the first node (12, Fig. 1) from the modified state to an owner state without updating memory (*e.g.*, Appellant's Specification, para. 0044), the second node (14, Fig. 1) being operative to receive the data from the first node (12, Fig. 1) and assign a shared state to an associated state of the data at the second node (14, Fig. 1) (*e.g.*, Appellant's Specification, paras. 0020-44).

Claim 12 recites:

A multi-processor network (100, Fig. 3), comprising:
memory (110, Fig. 3) for storing data (*E.g.*, Appellant's Specification, para.0058);
a first processor node (102, Fig. 3) having a first processor node cache line (116, Fig. 3) including the data, the first processor node (102, Fig. 3) cache line having an associated state, the associated state of the first processor node cache line being a modified state (*e.g.*, Appellant's Specification, paras. 0056, 0072);

a second processor node (104, Fig. 3) operative to provide a non-migratory source broadcast read request for the data, the second processor node (104, Fig. 3) having a second processor node cache line with an associated state (*e.g.*, Appellant's Specification, para. 0072);

the first processor node (102, Fig. 3) being programmed to respond to the non-migratory source broadcast read request of the second processor node (104, Fig. 3) by providing a shared data response to the second processor node (104, Fig. 3) and transitioning the associated state of the first processor node (102, Fig. 3) cache line from the modified state to an owner state without updating the memory (110, Fig. 3) with the data, the data being stored in the second processor node cache line, the associated state of the second processor

node (104, Fig. 3) cache line being assigned a shared state (*e.g.*, Appellant's Specification, para. 0072).

Claim 21 recites:

A computer system (10, Fig. 1) comprising:

a source processor (14, Fig. 1) having an associated source processor cache (24, Fig. 1) (*E.g.*, Appellant's Specification, para. 0023), the source processor (14, Fig. 1) being operative to issue a selected one of a non-migratory source broadcast (XREADN) request for data (*e.g.*, Appellant's Specification, para. 0040) and a migratory source broadcast (XREADM) request for data (*e.g.*, Appellant's Specification, para. 0044);

memory (16, Fig. 1) storing the data (*e.g.*, Appellant's Specification, para. 0021); and

a target processor (12, Fig. 1) having an associated target processor cache (22, Fig. 1) with a target processor cache line (116 in FIG. 3) that stores the data, the target processor cache line having an associated state, the associated state of the target processor cache line being a modified state (*e.g.*, Appellant's Specification, paras. 0040, 0061), the target processor (12, Fig. 1) being programmed to respond to the XREADN request by providing a shared data (S-DATA) response to the source processor (14, Fig. 1) and by transitioning the associated state of the target processor (12, Fig. 1) cache line from the modified state to an owner state without updating the memory (16, Fig. 1) (*e.g.*, Appellant's Specification, paras. 0044, 0072), the target processor (12, Fig. 1) being programmed to respond to the XREADM request by providing an ownership data (D-DATA) response to the source processor and by transitioning the associated state of the target processor (12, Fig. 1) cache line from the modified state to an invalid state without updating the memory (16, Fig. 1) (*e.g.*, Appellant's Specification, paras. 0040, 0070).

Claim 24 recites:

A system (10, Fig. 1, 50, Fig. 2, 100, Fig. 3, 150, Fig. 4, 170, Fig. 5, 190, Fig. 6) comprising means for broadcasting from a first node (14, Fig. 1, 60, Fig. 2, 104, Fig. 3, 192, Fig. 6), a non-migratory read (XREADN) request for data (*e.g.*, Appellant's Specification, paras. 0044, 0056, 0072, 0077);

means for providing the data from a second node (12, Fig. 1, 56, Fig. 2, 102, Fig. 3, 156, Fig. 4, 196, Fig. 6) to the first node (14, Fig. 1, 60, Fig. 2, 104, Fig. 3, 192, Fig. 6) in response to the XREADN request, a modified state being associated with the data at the second node (*e.g.*, Appellant's Specification, paras. 0044, 0056, 0072, 0077), a shared state being associated with the data at the first node (14, Fig. 1, 60, Fig. 2, 104, Fig. 3, 192, Fig. 6) in response to the first node (14, Fig. 1, 60, Fig. 2, 104, Fig. 3, 192, Fig. 6) receiving the data from the second node (12, Fig. 1, 56, Fig. 2, 102, Fig. 3, 156, Fig. 4, 196, Fig. 6) (*e.g.*, Appellant's Specification, paras. 0044, 0056, 0072, 0077); and

means for transitioning (122, Fig. 3) the modified state associated with the data at the second node (12, Fig. 1, 56, Fig. 2, 102, Fig. 3, 156, Fig. 4, 196, Fig. 6) to an owner state without updating memory (16, Fig. 1, 72, Fig. 2, 110, Fig. 3, 198, Fig. 6) of the system (10, Fig. 1, 50, Fig. 2, 100, Fig. 3, 150, Fig. 4, 170, Fig. 5, 190, Fig. 6) (*e.g.*, Appellant's Specification, paras. 0044, 0056, 0072, 0077).

Claim 28 recites:

A method comprising:

broadcasting (300, Fig. 8) a non-migratory request for data from a first node to other nodes of an associated system (*e.g.*, Appellant's Specification, para. 0079);

providing (310, Fig. 8) a shared copy of the data from a second node to the first node in response to the non-migratory request (*e.g.*, Appellant's Specification, para. 0079);

transitioning (320, Fig. 8) a state associated with the data at the second node from a modified state to an owner state in response to the non-migratory request (*e.g.*, Appellant's Specification, para. 0079); and

transitioning (330, Fig. 8) a state associated with the data at the first node to a shared state in response to receiving the shared copy of the data from the second node (*e.g.*, Appellant's Specification, para. 0079).

Claim 32 recites:

A computer system (10, Fig. 1) comprising a cache coherency protocol that is operative to permit migration of data to a cache (24, Fig. 1) associated with a source processor (14, Fig. 1) from a cache (22, Fig. 1) associated with a target processor (12, Fig. 1) when a migratory request is issued from the source processor (12, Fig. 1) (*e.g.*, Appellant's Specification, para. 0042), the protocol being further operative to prevent migration of the data to the cache (24, Fig. 1) associated with the source processor (14, Fig. 1) from the cache (22, Fig. 1) associated with the target processor (12, Fig. 1) when a non-migratory request is issued from the source processor (14, Fig. 1) (*e.g.*, Appellant's Specification, para. 0046).

VI. Grounds of Rejection to be Reviewed on Appeal

The Office Action raised the following grounds of rejection.

(1) Claims 1, 8-10, 12, 18-19, and 24-32 were rejected under 35 U.S.C. § 103(a) as being obvious over U.S. Patent Publication No. 2002/0129211 by Arimilli et al.

(“Arimilli”) in view of U.S. Patent No. 6,931,496 to Chen et al. (“Chen”).

(2) Claims 11 and 20-23 were rejected under 35 U.S.C. § 103(a) as being obvious over Arimilli in view of Chen and further in view of U.S. Patent No. 6,484,240 to Cypher et al. (“Cypher”).

Accordingly, Appellant hereby requests review of each of these grounds of rejection in the present appeal.

VII. Argument

(1) Claims 1, 8-10, 12, 18-19, and 24-32 are patentable over Arimilli and Chen:

Claims 1, 8-10, 12, 18-19, and 24-32 were rejected as being obvious over Arimilli in view of Chen. For at least the following reasons, this rejection should not be sustained.

Claim 1:

Claim 1 recites:

A system comprising:
a first node including data having an associated state, the associated state of the data at the first node being a modified state; and
a second node operative to provide a non-migratory source broadcast request for the data, the first node being operative in response to the non-migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an owner state without updating memory, the second node being operative to receive the data from the first node and assign a shared state to an associated state of the data at the second node.

(Emphasis added).

“The examiner bears the initial burden . . . of presenting a *prima facie* case of unpatentability.” *In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). A *prima facie* case of obviousness made under 35 U.S.C. § 103(a) requires a showing that all of the subject matter in the claim at issue would be obvious to one having ordinary skill in the art based on the teachings of the cited prior art. *See* M.P.E.P. § 2143.

The recent Office Action falls short of establishing a *prima facie* case of obviousness against the subject matter of claim 1 because it fails to provide sufficient evidence that the cited prior art references coupled with the knowledge of one having ordinary skill in the art would render all of the subject matter recited in claim 1 obvious at the time of the invention.

In the rejection of claim 1, it is contended in the Final Action that Arimilli teaches a second node, as recited in claim 1. (Action, p. 11) (citing to Arimilli, p. 7). Appellant

respectfully disagrees. In rejecting claim 1, the Final Action purports to separate interrelated features and alleges support for a first set of features are found in Arimilli and that support for a second set of features are found in a second reference. However, this approach fails to give proper consideration to the claim as a whole. The portions of claim 1 which the Final Action considers to be disclosed in Arimilli must be considered in view of the other claim language, such as including the state of the data at the first node and the particular type of request provided by the second node. The separate reliance on Arimilli appears to give insufficient weight and even ignores the interrelationships between the state of data at the nodes and the requests and responses recited in the system of claim 1.

Contrary to the contentions in the Final Action regarding claim 1, Arimilli fails to teach or suggest that the second node is operative to receive data from the first node and that the second node is operative to assign a shared state to an associated state of the data at the second node, as recited in claim 1. The Final Action relies entirely on the teachings at page 7 of Arimilli to support the rejection of claim 1. However, this portion of Arimilli discloses arbitrating between conflicting requests to modify a cache line that is held in a shared state and for protecting ownership of the cache line granted during such arbitration. (Arimilli, para. 0060). Nothing on page 7 of Arimilli (or anywhere else in Arimilli) teaches or suggests a second node receives data from a first node and assigning a shared state to an associated state of the data at the second node, such as recited in claim 1.

In Arimilli, if the coherency state associated with a target cache line has a state other than shared or invalid (including a modified state), a master 26 performs a store into cache array 24 without issuing a transaction on the system bus. (Arimilli, para. 0036). The master 26 in Arimilli prevents access to the target cache line by other agents 10 by means of

appropriate snoop responses until the store into the cache array is completed. (Arimilli, para. 0036).

In Arimilli, if the target cache line is invalid, the master 26 issues a read with intent to modify (RWITM) transaction on the system bus 12 to obtain a copy of the cache line from another agent 10 for modification. (Arimilli, para. 0036). If the master 26 receives a combined response awarding ownership of the target cache line for purposes of modifying the cache line (OWNER CR), the master 26 performs a store (of the target cache line) because all of the agents 10 caching the same (target) cache line have invalidated or will invalidate their respective copies of the cache line. (Arimilli, para. 0057 and Table III). If the agent 10 receives a combined response indicating that the master 26 has awarded ownership of the target cache line for purposes of modifying the cache line, but must perform clean up operations to maintain coherency (OWNER_CU CR), the master 26 issues high priority kill transactions on the system bus 12 before performing the store. (Arimilli, para. 0058, Table III and Fig. 3B). The cleanup operations, which are utilized to maintain coherency, are described in Arimilli in invalidating the cache line for all agents (including any previous owner), such that in contrast to claim 1 any prior owner who might have held the data in the modified state before the transaction now holds the data in the invalid (I) state. (Arimilli, paras. 0058, 0060).

Additionally, according to Arimilli, the agent that issues the read with intent to modify (RWITM) transaction seeks and is provided ownership, if the CDP grants ownership, such that the agent that issues the transaction (Dclaim or RWITM) does not store data in the shared state, but instead is awarded ownership. (Arimilli, paras. 0057, 0061).

The Final Action contends that paragraph 0029 of Arimilli also teaches certain features of claim 1; namely that a “second node operative to receive data from a first node and

that the second node is operative to assign a shared state to an associated state of the data at the shared node,” such as recited in claim 1. (Action, p. 4). Again, this contention fails to appreciate the particular interrelated set of circumstances recited in claim 1 that provide for and enable this result. Appellant respectfully submits that paragraph 0029 of Arimilli relates to an identification of a situation wherein a conflict can arise (e.g., between times t0 and t2) when a modifying transaction has been issued for a target cache line that is marked as shared in another agent’s cache. (Arimilli, para. 0029 and Fig. 2). Thus, Appellant respectfully submits that paragraph 0029 is unrelated to the subject matter recited in claim 1.

Furthermore, Appellant asserts that none of the situations taught or suggested by Arimilli describes a second node receives data from a first node having data in a modified state and assigning a shared state to the data at the second node, consistent with the system recited in claim 1. This is because in the approach disclosed in Arimilli, a target cache line is either: (a) already in the shared state (thus the data is not provided by another node) or (b) received from another agent 10, wherein, when the target cache line is received, the agent 10 is awarded (exclusive) ownership (as indicated by the OWNER CR or OWNER_CU CR) by the master 26, such that the target cache line is not assigned a shared state.

The Final Action concedes Arimilli is deficient as it “does not explicitly teach the first node being operative in response to the non-migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an owner state without updating system memory.” (Action, p. 4). Chen fails to make up for the deficiencies of Arimilli in this regard.

In contrast to the contentions in the Final Action, none of the sections of Chen cited in the Final Action (e.g., cols. 3, 5 and 6 and the abstract of Chen), nor Chen more generally, teach or suggest that a first node is operative to provide the data to the second node and

transition the associated state of the data at the first node from the modified state to an owner state, in response to a non-migratory source broadcast request provided by the second node as recited in claim 1. The systems and methods of Chen relate to and describe a directory based approach and not a source broadcast based system. (*See* Chen, col. 2, lines 14 to 25 and col. 4, lines 51-55).

Thus, the transactions for cached data in Chen, all go through a distributed shared memory (DSM) controller, such that no non-migratory source broadcast request would exist in the system of Chen. In Chen, the DSM controller maintains and manages states of the data (in an associated directory) for data in a level 3 (L3) cache based on the commands issued by the respective processors. (Chen, col. 2, lines 20-25). Since no non-migratory source broadcast request exists in Chen, there is no basis to conclude that the teachings of Chen would be operative to provide data to the second node in response to such a non-migratory broadcast request. Therefore, Arimilli taken in view of Chen fails to teach or suggest that a first node is operative, in response to a non-migratory source broadcast request, to provide data to a second node, as recited in claim 1.

Furthermore, Arimilli taken in view of Chen fails to teach or suggest that, in response to the non-migratory source broadcast request, a first node transitions the state of data at such node from a modified state to an owner state without updating system memory, as recited in claim 1. In rejecting claim 1, the Final Action contends that Chen discloses this element of claim 1. (Action, p. 14). Appellant respectfully disagrees. Chen discloses seven possible states for each of the L3 cache lines: CLEAN, FRESH, DIRTY-ONLY, DIRTY-SHARED, VOID, IDLE. (Chen col. 4, lines 60-65). Of the L3 cache states disclosed in Chen, only the DIRTY-ONLY and DIRTY-SHARED states indicate that data has been modified. (Chen, col. 5, lines 6 to 14).

Chen fails to teach or suggest that any command, (even a BRIL command, which is a command issued to read an exclusive copy of the specific data, *see* Chen Col. 5, Lines 32-33), which enables a second node to receive data from a L3 cache having the data in either of the DIRTY-ONLY and DIRTY-SHARED states and then sets the state of all nodes to a shared state. (Chen, col. 7, lines 48-53). As discussed above, the approach in Chen employs directories in a DSM controller to manage all requests and the states are not associated with data at more than one node, but instead the states apply to the data in the L3 cache of the local node.

Accordingly, Chen fails to teach or suggest transitioning from a modified state to an owner state in response to the non-migratory source broadcast request, as recited in claim 1, since in Chen, no command is taught or suggested that enables a second node to receive data from a node with data in either DIRTY-ONLY or DIRTY-SHARED states and then set the state to a shared state.

Further still, the Court of Customs and Patent Appeals has held that if a proposed modification or combination of the prior art would change the principle of operation of the prior art invention being modified, then the teachings of the references are not sufficient to render the claims *prima facie* obvious. *In re Ratti*, 270 F.2d 810, 813, 270 F.2d 810 (C.C.P.A. 1959). Appellant asserts that one of ordinary skill in the art would not combine and modify the teachings of Chen comprising a DSM system with the system of Arimilli comprising a central memory system because the person of ordinary skill would recognize that the transactions required to maintain coherency in a DSM system would adversely affect the principal operation of a central memory system. (Arimilli's System Memory 12; *see* Chen Fig. 1 and Arimilli Fig 1). Moreover, the combined teachings of Arimilli and Chen still fail

to teach or suggest all features recited in claim 1, such that claim 1 would not be obvious to one of ordinary skill in the art.

Moreover, the Final Action contends that one of ordinary skill in the art would see to modify Arimilli to include the transitioning of states from a modified state to an owner state without updating system memory because this would have provided a data maintenance method in a distributed shared memory system to efficiently solve the access deadlock problem as taught by Chen. (Action, pp. 4-5). However, the approach in Arimilli is directed to a specific approach in which a coherency decision point (CDP) is utilized as a mechanism to arbitrate between conflict requests to modify data. (Arimilli, paras 0011, 0030). It is respectfully submitted, that the approach taught in Chen is not consistent with and would likely compromise the intent of the system of Arimilli. The contention to combine the teachings of Arimilli and Chen appear to be unsupported by the record and appear to be based improperly on speculation or the Examiner's subjective belief. *In re Lee*, 277 F.3d 1338, 61 U.S.P.Q.2d 1430 (Fed. Cir 2002).

Appellant notes that the present rejection relies on an assertion of obviousness based on the combination of prior art elements according to known methods to yield predictable results. According to the M.P.E.P., "[t]o reject a claim based on this rationale, Office personnel must resolve the . . . factual inquiries" of *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1 (1966). M.P.E.P. § 2143.

Under the analysis required by *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1 (1966), to support a rejection under § 103, the scope and content of the prior art must first be determined, followed by an assessment of the differences between the prior art and the claim at issue in view of the ordinary skill in the art. The Supreme Court has recently reaffirmed that the *Graham* factors "continue to define the inquiry that controls" obviousness rejections

under § 103. *KSR Int'l v. Teleflex Inc.*, 550 U.S. 398, 407 (2007). In the present case, the scope and content of the prior art, as evidenced by Arimilli and Chen, did not include the claimed subject matter, particularly “a second node operative to provide a non-migratory source broadcast request for the data, the first node being operative in response to the non-migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an owner state without updating memory, the second node being operative to receive the data from the first node and assign a shared state to an associated state of the data at the second node.” (Claim 1).

The differences between the cited prior art and the claim 1 are significant because the claimed subject matter provides features and advantages not known or available in the cited prior art. Consequently, the cited prior art will not support a rejection of claim 1 under 35 U.S.C. § 103 and *Graham*. For at least these reasons, the rejection of claim 1 and its dependent claims based on Arimilli and Chen should not be sustained.

Claim 8:

Claim 8 depends from claim 1 and is patentable over the cited prior art for at least the same reasons given above in favor of the patentability of claim 1. Additionally, claim 8 recites subject matter that is further patentable over the prior art of record.

Specifically, claim 8 recites “wherein the first node is operative in response to the non-migratory source broadcast request to provide a shared data response to the second node.” In the rejection of claim 8, it is contended in the Final Action that Chen discloses the elements of claim 8. (Action, p. 5). Appellant respectfully disagrees. Chen fails to teach or suggest that a first node is operative, in response to a non-migratory source broadcast request, to provide a shared data response to a second node, as recited in claim 8. In fact, Chen fails

to teach or suggest a source broadcast request - it employs a DSM controller with a directory for managing data in an L3 cache - such that the system disclosed in Chen is not operative to respond to non-migratory source broadcast requests, in contrast to the first node recited in claim 8. Accordingly, Arimilli taken in view of Chen fails to teach or suggest the system recited in claim 8 since Chen fails to teach or suggest a source broadcast system. Therefore, for at least these additional reasons, the rejection of claim 8 should not be sustained.

Claim 12:

Claim 12 recites:

A multi-processor network comprising:

memory for storing data;

a first processor node having a first processor node cache line including the data, the first processor node cache line having an associated state, the associated state of the first processor node cache line being a modified state; and

a second processor node operative to provide a non-migratory source broadcast read request for the data, the second processor node having a second processor node cache line with an associated state;

the first processor node being programmed to respond to the non-migratory source broadcast read request of the second processor node by providing a shared data response to the second processor node and transitioning the associated state of the first processor node cache line from the modified state to an owner state without updating the memory with the data, the data being stored in the second processor node cache line, the associated state of the second processor node cache line being assigned a shared state.

(Emphasis added).

Arimilli taken in view of Chen fails to make claim 12 obvious. Arimilli taken in view of Chen fails to teach or suggest data (requested by a non-migratory source broadcast read request) is stored in a second processor node cache line, wherein an associated state of the second processor node cache line is assigned a shared state, as recited in claim 12. In contrast to claim 12, Arimilli teaches that a target cache line is already in a shared state (such that data is not provided by another node) or the target cache line is received from another agent 10,

wherein the ownership of the target cache line is awarded by a master 26 to a requesting agent 10, such that the target cache line would not be assigned a shared state. (Arimilli, Figs. 3A and 3B). Therefore, Arimilli taken in view of Chen fails to teach or suggest the second processor node operating in the manner recited in claim 12.

Moreover, Arimilli taken in view of Chen fails to teach or suggest a first processor node that is programmed to respond to a non-migratory source broadcast read request of the second processor node by providing a shared data response to the second processor node, as recited in claim 12. The Final Action contends that Chen discloses this element of claim 12 (Action, p. 4, in rejecting claim 1), but in contrast to claim 1, Chen teaches a directory based cache coherency system for an L3 cache that is managed by the a DSM controller. (Chen, col. 2, lines 14 to 24 and col. 4, lines 51-55). The set of requests and responses taught in Chen relate to requests issued to and responses from the DSM controller for data in the L3 cache. Chen also fails to teach the use of any non-migratory source broadcast read request issued by a source processor node, as recited in claim 12.

Accordingly, Arimilli taken in view of Chen cannot teach or suggest that any response (from a first node) to a non-migratory source broadcast read request can be provided to a second node, since Chen fails to teach or suggest a source broadcast based system. Additionally, Arimilli taken in view of Chen fails to teach or suggest transitioning an associated state of the first processor from a modified state to an owner state without updating the memory, as recited in claim 12. In Chen, which the Final Action contends discloses this element of claim 12 (Action, pp. 4-6), only the DIRTY-ONLY and DIRTY-SHARED states indicate that data in an L3 cache has been modified. (Chen, col. 5, lines 6 to 14).

Chen fails to teach or suggest any command that enables a second processor node to receive data from a processor node with the data in either of the DIRTY-ONLY and DIRTY-

SHARED states and then set the state to a shared state. Instead, Chen teaches that the cache states and transitions thereof relate to the cache states of the L3 cache (Chen, col. 4, line 60 through col. 5, line 21), which are not cache states for data stored in the cache line of a processor node, as recited in claim 12. Since Chen fails to teach or suggest a command that enables a second processor node to receive data from a processor node with the data in either of the DIRTY-ONLY and DIRTY-SHARED states and then set the state to a shared state, Arimilli taken in view of Chen also fails to teach or suggest transitioning the associated state of the first processor from a modified state to an owner state without updating the memory, as recited in claim 12.

Furthermore, combining and modifying the teachings of Arimilli and Chen in the manner suggested in the Final action would change a basic principle of operation of Arimilli's central memory system for the reasons stated above with respect to claim 1. Additionally, there is not proper motivation to combine Arimilli and Chen for the reasons discussed above with respect to claim 1. Motivation to modify Arimilli based on teachings in Chen to provide the system of claim 12 is further lacking since Chen relates to management of an L3 cache by DSM controller whereas claim 12 recites first and second processor nodes having cache lines and associated cache states.

Again, under the analysis required by *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1 (1966), to support a rejection under § 103, the scope and content of the prior art must first be determined, followed by an assessment of the differences between the prior art and the claim at issue in view of the ordinary skill in the art. The Supreme Court has recently reaffirmed that the *Graham* factors “continue to define the inquiry that controls” obviousness rejections under § 103. *KSR Int’l v. Teleflex Inc.*, 550 U.S. 398, 407 (2007). In the present case, the scope and content of the prior art, as evidenced by Arimilli and Chen, did not

include the claimed subject matter, particularly “a second processor node operative to provide a non-migratory source broadcast read request for the data, the second processor node having a second processor node cache line with an associated state;” and “the first processor node being programmed to respond to the non-migratory source broadcast read request of the second processor node by providing a shared data response to the second processor node and transitioning the associated state of the first processor node cache line from the modified state to an owner state without updating the memory with the data, the data being stored in the second processor node cache line, the associated state of the second processor node cache line being assigned a shared state.” (Claim 12)

The differences between the cited prior art and the claim 12 are significant because the claimed subject matter provides features and advantages not known or available in the cited prior art. Consequently, the cited prior art will not support a rejection of claim 12 under 35 U.S.C. § 103 and *Graham*. For at least these reasons, the rejection of claim 12 and its dependent claims should not be sustained.

Claims 10 and 19:

Claim 10 depends from claim 1 and Claim 19 depends from claim 12. Accordingly, claims 10 and 19 are patentable over the cited prior art for at least the same reasons given above in favor of the patentability of claims 1 and 12, respectively. Additionally, claim 18 recites subject matter that is further patentable over the prior art of record.

Specifically, claim 10 recites “at least one other node that provides a non-data response to the second node in response to the non-migratory source broadcast request from the second node, the non-data response indicating that the at least one other node does not have a valid copy of the data requested by the second node.” Similarly, claim 19 recites “at

least one other processor node that provides a non-data response to the second processor node in response to the non-migratory source broadcast read request from the second processor node, the non-data response indicating that the at least one other processor node does not have a valid copy of the data requested by the second processor node.”

In rejecting claim 19, the Action relies solely on the basis for the rejection of claim 10. (Action, p. 6). In rejecting claim 10, it is contended in the Final Action that col. 7, lines 55 to 61 of Chen discloses the elements recited in claim 10. (Action, p. 6). Appellant respectfully disagrees. The cited section of Chen is unrelated to a non-data response that is provided to a second (processor) node, in contrast to the system recited in claims 10 and 19. Moreover, Chen fails to teach or suggest a relationship between three nodes (or processor nodes, as recited in claim 19): first, second and at least one other node - and also fails to teach or suggest a system with source broadcast requests, consistent with what is recited in claims 10 and 19.

Specifically, by virtue of claims 10 and 19's dependence from claims 1 and 12, claims 10 and 19 recite a second node that broadcasts a non-migratory source request for data, a first node that provides the data to the second node, and at least one other node (e.g., a third node) that provides a non-data response to the second node in response to the non-migratory source broadcast request from the second node, the non-data response indicating that the at least one other node does not have a valid copy of the data requested by the second node. Additionally, in claim 19, the nodes are specifically recited as being processor nodes having cache lines.

In contrast to what appears is being suggested in the Final Action, the commands LRIL or LIL are not responses to non-migratory source broadcast requests, but instead, LRIL and LIL are commands issued to read data. (Chen, col. 5, lines 48-50). Significantly, no responses of any kind are disclosed as being issued by any nodes in the cited section of Chen.

(E.g., Chen, col. 7, lines 55 to 61). Therefore, Appellant respectfully submits that Arimilli taken in view of Chen fails to make claims 10 and 19 obvious since Arimilli taken in view of Chen fails to teach or suggest the interrelationship between the first, second and third node (e.g., the at least one other node) recited in claims 10 and 19. Additionally, in claim 19, the non-data response is provided by a processor node.

For at least these additional reasons, the rejection of claims 10 and 19 should not be sustained.

Claim 24:

Claim 24 recites:

A system comprising:
means for broadcasting from a first node a non-migratory read (XREADN) request for data;
means for providing the *data from a second node to the first node in response to the XREADN request, a modified state being associated with the data at the second node, a shared state being associated with the data at the first node in response to the first node receiving the data from the second node*; and
means for *transitioning the modified state associated with the data at the second node to an owner state without updating memory of the system*.
(Emphasis added).

Arimilli taken in view of Chen fails to make claim 24 obvious. In the Final Rejection, the rationale for the rejection of claim 1 is relied on as the sole basis for the rejection of claim 24. (Action, p. 6). Arimilli taken in view of Chen fails to teach or suggest data (requested by a non-migratory source broadcast read request (XREADN)) is provided to a first node, wherein a shared state is associated with the data at the first node in response to the first node receiving the data from a second node, as recited in claim 24. In contrast to the system recited in claim 24, Arimilli teaches that a target cache line is already in a shared state (such that data is not provided by another node) or the target cache line is received from another

agent 10, wherein the ownership of the target cache line is awarded by a master 26 to a requesting agent 10, such that the target cache line would not be assigned a shared state. (Arimilli, Figs. 3A and 3B). Therefore, Arimilli taken in view of Chen fails to teach or suggest the means for providing recited in claim 24.

Moreover, Arimilli taken in view of Chen fails to teach or suggest means for providing data from a second node to the first node in response to a XREADN request, as recited in claim 24. In Chen, which the Final Action contends discloses this element of claim 24 (Action, p. 4, in rejecting claim 1), a directory based cache system is disclosed (Chen, col. 2, lines 14 to 24 and col. 4, lines 51-55) in which a DSM controller manages states of cache lines in an L3 cache. In contrast to claim 24, Chen is silent on the employment of any source broadcast based system. Accordingly, Arimilli taken in view of Chen cannot teach or suggest that any response (from a first node) to an XREADN can be provided to a second node since Chen fails to teach or suggest a source based broadcast system. Therefore, Arimilli taken in view of Chen fails to teach or suggest the means for providing recited in claim 24.

Additionally, Arimilli taken in view of Chen fails to teach or suggest means for transitioning a modified state associated with data at a second node to an owner state without updating the memory, as recited in claim 24. In Chen, which the Final Action contends discloses this element of claim 24 (Action, p. 4), only the DIRTY-ONLY and DIRTY-SHARED states indicate that data in an L3 cache has been modified. (Chen Col. 5, Lines 6 to 14). Chen fails to teach or suggest any non-migratory command that would enable a second node to receive data from a node with the data in either of the DIRTY-ONLY and DIRTY-SHARED states and then set the state to a shared state upon receipt, while also transitioning the data in either of the DIRTY-ONLY and DIRTY-SHARED states to an owner state, as recited in claim 24.

Furthermore, combining and modifying the teachings of Arimilli and Chen in the manner suggested in the Final Action would change a basic principle of operation of Arimilli's central memory system (System Memory 24) for the reasons discussed above with respect to claim 1. Additionally, it would not be reasonable for one of ordinary skill in the art to modify Arimilli in view of the teachings of Chen to provide the system of claim 24 for the reasons discussed above with respect to claim 1.

Again, under the analysis required by *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1 (1966), to support a rejection under § 103, the scope and content of the prior art must first be determined, followed by an assessment of the differences between the prior art and the claim at issue in view of the ordinary skill in the art. The Supreme Court has recently reaffirmed that the *Graham* factors “continue to define the inquiry that controls” obviousness rejections under § 103. *KSR Int’l v. Teleflex Inc.*, 550 U.S. 398, 407 (2007). In the present case, the scope and content of the prior art, as evidenced by Arimilli and Chen, did not include the claimed subject matter, particularly “means for providing the data from a second node to the first node in response to the XREADN request, a modified state being associated with the data at the second node, a shared state being associated with the data at the first node in response to the first node receiving the data from the second node;” and “means for transitioning the modified state associated with the data at the second node to an owner state without updating memory of the system.” (Claim 24).

The differences between the cited prior art and the claim 24 are significant because the claimed subject matter provides features and advantages not known or available in the cited prior art. Consequently, the cited prior art will not support a rejection of claim 24 under 35 U.S.C. § 103 and *Graham*. For at least these reasons, the rejection of claim 24 and its dependent claims should not be sustained.

Claim 25:

Claim 25 depends from claim 24 and is patentable over the cited prior art for at least the same reasons given above in favor of the patentability of claim 24. Additionally, claim 25 recites subject matter that is further patentable over the prior art.

Appellant notes that claim 25 recites similar subject matter to dependent claim 2. The Office has conceded that claim 2 recites allowable subject matter. (Action, p. 10). Consequently, Appellant respectfully submits that claim 25 is patentable over Arimilli and Chen for at least the same reasons that claim 2 is patentable over Arimilli and Chen.

Claims 26-27:

Claims 26-27 depend from claim 25 and are patentable over the cited prior art for at least the same reasons given above in favor of the patentability of claims 24 and 25. Additionally, claims 26-27 recite subject matter that is further patentable over the prior art.

Specifically, claim 26 recites “means for selecting one of the XREADM request and XREADN request to broadcast from the first node.” Claim 27 recites the similar subject matter of “predictively selecting one of the XREADM request and XREADN request to broadcast from the first node.” The Action asserts that Figs. 4-5 of Chen render claims 26-27 obvious. Appellant respectfully disagrees. As will be readily apparent from a careful examination of Chen, Figs. 4-5 of Chen do not teach or suggest anywhere the selection between an XREADM request and an XREADN request to broadcast from the first node.

According to the Supreme Court, the Examiner is required to provide an explicit analysis as to how the cited prior art teaches or suggests all the features of a claim. *KSR Int’l Co. v. Teleflex, Inc.*, 550 U.S. 398, 418 (2007) (citing to *In re Kahn*, 441 F.3d 977, 988 (Fed.

Cir. 2006)). “To facilitate review, this [the Examiner’s] analysis should be made explicit.”

Id. As demonstrated above, the Examiner has failed to meet this burden with regard to claim 26-27, by failing to identify which portions of Figs. 4-5 in Chen the Examiner believes to teach or suggest the subject matter of claims 26-27. Moreover, Chen does not appear to teach or suggest the subject matter of claims 26-27 anywhere. Therefore, under the standard of *KSR*, no *prima facie* case of obviousness has been made as to claims 26-27. For at least this additional reason, the rejection of claims 26-27 should not be sustained.

Claim 28:

Claim 28 recites:

A method comprising:
broadcasting a non-migratory request for data from a first node to other nodes of an associated system;
providing a shared copy of the data from a second node to the first node in response to the non-migratory request;
transitioning a state associated with the data at the second node from a modified state to an owner state in response to the non-migratory request; and
transitioning a state associated with the data at the first node to a shared state in response to receiving the shared copy of the data from the second node.
(Emphasis added).

Arimilli and Chen also fail to render the method of claim 28 obvious. Specifically, as demonstrated above with regard to claim 1, Arimilli and Chen fail to teach or suggest “providing a shared copy of the data from a second node to the first node in response to the non-migratory request;” “transitioning a state associated with the data at the second node from a modified state to an owner state in response to the non-migratory request;” and “transitioning a state associated with the data at the first node to a shared state in response to receiving the shared copy of the data from the second node.” (Claim 28).

Again, under the analysis required by *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1 (1966), to support a rejection under § 103, the scope and content of the prior art must first be determined, followed by an assessment of the differences between the prior art and the claim at issue in view of the ordinary skill in the art. The Supreme Court has recently reaffirmed that the *Graham* factors “continue to define the inquiry that controls” obviousness rejections under § 103. *KSR Int’l v. Teleflex Inc.*, 550 U.S. 398, 407 (2007). In the present case, the scope and content of the prior art, as evidenced by Arimilli and Chen, did not include the claimed subject matter, particularly “providing a shared copy of the data from a second node to the first node in response to the non-migratory request;” “transitioning a state associated with the data at the second node from a modified state to an owner state in response to the non-migratory request;” and “transitioning a state associated with the data at the first node to a shared state in response to receiving the shared copy of the data from the second node.” (Claim 28).

The differences between the cited prior art and the claim 28 are significant because the claimed subject matter provides features and advantages not known or available in the cited prior art. Consequently, the cited prior art will not support a rejection of claim 28 under 35 U.S.C. § 103 and *Graham*. For at least these reasons, the rejection of claim 28 and its dependent claims should not be sustained.

Claim 29:

Claim 29 depends from claim 28 and is patentable over the cited prior art for at least the same reasons given above in favor of the patentability of claim 28. Additionally, claim 29 recites subject matter that is further patentable over the prior art.

Appellant notes that claim 29 recites similar subject matter to dependent claim 2. The Office has conceded that claim 2 recites allowable subject matter. (Action, p. 10). Consequently, Appellant respectfully submits that claim 29 is patentable over Arimilli and Chen for at least the same reasons that claim 2 is patentable over Arimilli and Chen.

Claims 30-31:

Claims 30-31 depends from claim 29 and are patentable over the cited prior art for at least the same reasons given above in favor of the patentability of claims 28 and 29. Additionally, claims 30-31 recite subject matter that is further patentable over the prior art.

Specifically, claim 30 recites “means for selecting one of the migratory request and non-migratory request to broadcast from the first node.” Claim 27 recites the similar subject matter of “predictively selecting one of the migratory request and the non-migratory request to broadcast from the first node.” The Action asserts that Figs. 4-5 of Chen render claims 29-30 obvious. Appellant respectfully disagrees. As will be readily apparent from a careful examination of Chen, Figs. 4-5 of Chen do not teach or suggest anywhere the selection between a non-migratory request and a migratory request to rebroadcast from the first node.

According to the Supreme Court, the Examiner is required to provide an explicit analysis as to how the cited prior art teaches or suggests all the features of a claim. *KSR Int’l Co. v. Teleflex, Inc.*, 550 U.S. 398, 418 (2007) (citing to *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006)). “To facilitate review, this [the Examiner’s] analysis should be made explicit.” *Id.* As demonstrated above, the Examiner has failed to meet this burden with regard to claims 30-31, by failing to identify which portions of Figs. 4-5 in Chen the Examiner believes to teach or suggest the subject matter of claims 30-31. Moreover, Chen does not appear to teach or suggest the subject matter of claims 30-31 anywhere. Therefore, under the standard of

KSR, no *prima facie* case of obviousness has been made as to claims 30-31. For at least this additional reason, the rejection of claims 30-31 should not be sustained.

Claim 32:

Claim 32 recites:

A computer system comprising a cache coherency protocol that is operative to permit migration of data to a cache associated with a source processor from a cache associated with a target processor *when a migratory request is issued from the source processor, the protocol being further operative to prevent migration of the data to the cache associated with the source processor from the cache associated with the target processor when a non-migratory request is issued from the source processor.*

(Emphasis added).

Appellant respectfully submits that the Final Rejection of claim 32 fails to establish a *prima facie* case of obviousness with respect to claim 32. Specifically, as demonstrated above with regard to claim 1, Arimilli and Chen fail to teach or suggest “permit[ting] migration of data to a cache associated with a source processor from a cache associated with a target processor when a migratory request is issued from the source processor, the protocol being further operative to prevent migration of the data to the cache associated with the source processor from the cache associated with the target processor when a non-migratory request is issued from the source processor.” (Claim 32).

Again, under the analysis required by *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1 (1966), to support a rejection under § 103, the scope and content of the prior art must first be determined, followed by an assessment of the differences between the prior art and the claim at issue in view of the ordinary skill in the art. The Supreme Court has recently reaffirmed that the *Graham* factors “continue to define the inquiry that controls” obviousness rejections under § 103. *KSR Int’l v. Teleflex Inc.*, 550 U.S. 398, 407 (2007). In the present case, the scope and content of the prior art, as evidenced by Arimilli and Chen, did not

include the claimed subject matter, particularly “permit[ting] migration of data to a cache associated with a source processor from a cache associated with a target processor when a migratory request is issued from the source processor, the protocol being further operative to prevent migration of the data to the cache associated with the source processor from the cache associated with the target processor when a non-migratory request is issued from the source processor.” (Claim 32).

The differences between the cited prior art and the claim 32 are significant because the claimed subject matter provides features and advantages not known or available in the cited prior art. Consequently, the cited prior art will not support a rejection of claim 32 under 35 U.S.C. § 103 and *Graham*. For at least these reasons, the rejection of claim 32 and its dependent claims should not be sustained.

(2) Claims 11 and 20-23 are patentable over Arimilli, Chen, and Cypher:

Claims 11 and 20-23 were rejected as obvious over Arimilli in view of Chen and in further view of Cypher. For at least the following reasons, this rejection should not be sustained.

Claims 11 and 20:

Claims 11 and 20 are dependent from claims 1 and 12, respectively, and are patentable over the prior art for at least the same reasons given above in favor of the patentability of claims 1 and 12, respectively, because if an independent claim is nonobvious under 35 U.S.C. § 103, then any claim depending therefrom is nonobvious. *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988); *see also* M.P.E.P. § 2143.03. Consequently, for at least this reason, the rejection of claims 11 and 20 should not be sustained.

Claim 21:

Claim 21 recites:

A computer system comprising:

a source processor having an associated source processor cache, the source processor being operative to issue a selected one of a non-migratory source broadcast (XREADN) request for data and a migratory source broadcast (XREADM) request for data;

memory storing the data; and

a target processor having an associated target processor cache with a target processor cache line that stores the data, the target processor cache line having an associated state, the associated state of the target processor cache line being a modified state, *the target processor being programmed to respond to the XREADN request by providing a shared data (S-DATA) response to the source processor and by transitioning the associated state of the target processor cache line from the modified state to an owner state without updating the memory, the target processor being programmed to respond to the XREADM request by providing an ownership data (D-DATA) response to the source processor and by transitioning the associated state of the target processor cache line from the modified state to an invalid state without updating the memory.*

(Emphasis added).

Claim 21 is not made obvious by Arimilli taken in view of Chen and in further view of Cypher. In rejecting claim 21, the Action contends that claim 21 is unpatentable over Arimilli, in view of Chen and further in view of Cypher by incorporating the rationale used in the rejection of claims 1 and 11. (Action, p. 9). Appellant respectfully disagrees with this contention. Appellant respectfully submits that the Final Rejection of claim 21 fails to establish a *prima facie* case of obviousness with respect to claim 21. Appellant's representative respectfully submits that claim 21 recites elements similar to claim 2, which depends from claim 1 and was indicated as being allowable. Accordingly, Appellant's representative respectfully submits that claim 21 is patentable for reasons similar to claim 2.

Claims 22-23:

Claims 22 and 23 are dependent from claim 21, and are patentable over the prior art for at least the same reasons given above in favor of the patentability of claims 21, because if an independent claim is nonobvious under 35 U.S.C. § 103, then any claim depending therefrom is nonobvious. *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988); *see also* M.P.E.P. § 2143.03. Consequently, for at least this reason, the rejection of claims 22-23 should not be sustained.

In view of the foregoing, it is submitted that the final rejection of the pending claims is improper and should not be sustained. Therefore, a reversal of the Rejection of December 10, 2009 is respectfully requested.

Respectfully submitted,

DATE: February 10, 2010

/Steven L. Nichols/
Steven L. Nichols
Registration No. 40,326

Steven L. Nichols, Esq.
Managing Partner, Utah Office
Rader Fishman & Grauer PLLC
River Park Corporate Center One
10653 S. River Front Parkway, Suite 150
South Jordan, Utah 84095
(801) 572-8066
(801) 572-7666 (fax)

VIII. CLAIMS APPENDIX

1. (original) A system comprising:
 - a first node including data having an associated state, the associated state of the data at the first node being a modified state; and
 - a second node operative to provide a non-migratory source broadcast request for the data, the first node being operative in response to the non-migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an owner state without updating memory, the second node being operative to receive the data from the first node and assign a shared state to an associated state of the data at the second node.
2. (original) The system of claim 1, wherein the second node is further operative to provide a migratory source broadcast request for the data, the first node being operative in response to the migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an invalid state without updating memory, the second node being operative to receive the data from the first node and assign the associated state of the data at the second node to a dirty state.
3. (original) The system of claim 2, wherein the first node is operative in response to the migratory source broadcast request to provide an ownership data response to the second node.

4. (original) The system of claim 2, wherein the associated state of the data at the second node being the dirty state makes the data at the second node available for migration to other nodes.
5. (original) The system of claim 2, wherein the second node is further operative to write the data to the second node and transition the associated state of the data at the second node from the dirty state to the modified state, the associated state of the data at the second node being the modified state making the data at the second node available for migration to other nodes providing a migratory source broadcast request for the data.
6. (original) The system of claim 2, wherein the second node is programmed with instructions to selectively invoke one of the non-migratory source broadcast request and the migratory source broadcast request.
7. (original) The system of claim 2, wherein the second node is programmed with instructions that provide a predictive selection to invoke one of the non-migratory source broadcast request and the migratory source broadcast request.
8. (original) The system of claim 1, wherein the first node is operative in response to the non-migratory source broadcast request to provide a shared data response to the second node.
9. (original) The system of claim 1, wherein further migration of the data from the second node is precluded when the associated state of the data at the second node is the shared state.

10. (original) The system of claim 1, further comprising at least one other node that provides a non-data response to the second node in response to the non-migratory source broadcast request from the second node, the non-data response indicating that the at least one other node does not have a valid copy of the data requested by the second node.

11. (original) The system of claim 1, wherein the first node defines a first processor and the second node defines a second processor, the first and second processors each having an associated cache that comprises a plurality of cache lines, each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the respective cache line, the first and second processors being capable of communicating with each other and with a system memory via an interconnect, the system further comprising a first cache controller associated with the first processor and a second cache controller associated with the second processor, the first cache controller being operative to manage data requests and responses for the associated cache of the first processor, the first cache controller effecting state transitions associated with the data in the associated cache of the first processor based on the data requests and responses for the associated cache of the first processor, the second cache controller being operative to manage data requests and responses for the associated cache of the second processor, the second cache controller effecting state transitions associated with the data in the associated cache of the second processor based on the data requests and responses for the associated cache of the second processor.

12. (original) A multi-processor network comprising:

memory for storing data;

a first processor node having a first processor node cache line including the data, the first processor node cache line having an associated state, the associated state of the first processor node cache line being a modified state; and

a second processor node operative to provide a non-migratory source broadcast read request for the data, the second processor node having a second processor node cache line with an associated state;

the first processor node being programmed to respond to the non-migratory source broadcast read request of the second processor node by providing a shared data response to the second processor node and transitioning the associated state of the first processor node cache line from the modified state to an owner state without updating the memory with the data, the data being stored in the second processor node cache line, the associated state of the second processor node cache line being assigned a shared state.

13. (original) The network of claim 12, wherein the second processor node is further operative to provide a migratory source broadcast read request for the data, the first processor node being programmed to respond to the migratory source broadcast read request of the second processor node by providing an ownership data response to the second processor node and transitioning the associated state of the first processor node cache line from the modified state to an invalid state without updating the memory with the data, the data from the ownership data response being stored in the second processor node cache line and the state associated with the second processor node cache line being assigned a dirty state.

14. (original) The network of claim 13, wherein the data stored in the second processor node cache line assigned the dirty state is available for migration to other nodes.

15. (original) The network of claim 13, wherein the second processor node is further operative to write the data stored in the second processor node cache line assigned the dirty state and transition the state associated with the second processor node cache line from the dirty state to the modified state, the data stored in the second processor node cache line assigned the modified state being available for migration to other nodes providing a migratory source broadcast read request for the data.

16. (original) The network of claim 13, wherein the second processor node is programmed with instructions to selectively invoke one of the non-migratory source broadcast read request and the migratory source broadcast read request to obtain the data.

17. (original) The network of claim 13, wherein the second processor node is programmed with instructions to predictively select one of the non-migratory source broadcast read request and the migratory source broadcast read request to obtain the data.

18. (original) The network of claim 12, wherein further migration of the data from the second processor node is precluded when in the shared state.

19. (original) The network of claim 12, further comprising at least one other processor node that provides a non-data response to the second processor node in response to the non-migratory source broadcast read request from the second processor node, the non-data

response indicating that the at least one other processor node does not have a valid copy of the data requested by the second processor node.

20. (original) The network of claim 12, wherein the first and second processor nodes each have an associated cache that comprises a plurality of cache lines, each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the respective cache line, the first and second processor nodes being capable of communicating with each other and with the memory via an interconnect, the system further comprising a first cache controller associated with the first processor node and a second cache controller associated with the second processor node, the first cache controller being operative to manage data requests and responses for the associated cache of the first processor node, the first cache controller effecting state transitions associated with the data in the cache of the first processor node based on the data requests and responses for the associated cache of the first processor node, the second cache controller being operative to manage data requests and responses for the associated cache of the second processor node, the second cache controller effecting state transitions associated with the data in the associated cache of the second processor node based on the data requests and responses for the associated cache of the second processor node.

21. (original) A computer system comprising:

a source processor having an associated source processor cache, the source processor being operative to issue a selected one of a non-migratory source broadcast (XREADN) request for data and a migratory source broadcast (XREADM) request for data;

memory storing the data; and

a target processor having an associated target processor cache with a target processor cache line that stores the data, the target processor cache line having an associated state, the associated state of the target processor cache line being a modified state, the target processor being programmed to respond to the XREADN request by providing a shared data (S-DATA) response to the source processor and by transitioning the associated state of the target processor cache line from the modified state to an owner state without updating the memory, the target processor being programmed to respond to the XREADM request by providing an ownership data (D-DATA) response to the source processor and by transitioning the associated state of the target processor cache line from the modified state to an invalid state without updating the memory.

22. (original) The computer system of claim 21, wherein the source processor further comprises an associated source processor cache having a source processor cache line for storing the data, the source processor cache line having an associated state, the source processor storing the data in the source processor cache line and assigning a shared state to the associated state of the source processor cache line in response to receiving the S-DATA response from the target processor.

23. (original) The computer system of claim 21, wherein the source processor further comprises an associated source processor cache having a source processor cache line for storing the data, the source processor cache line having an associated state, the source processor storing the data in the source processor cache line and assigning a dirty state to the

associated state of the source processor cache line in response to receiving the D-DATA response from the target processor.

24. (original) A system comprising:

means for broadcasting from a first node a non-migratory read (XREADN) request for data;

means for providing the data from a second node to the first node in response to the XREADN request, a modified state being associated with the data at the second node, a shared state being associated with the data at the first node in response to the first node receiving the data from the second node; and

means for transitioning the modified state associated with the data at the second node to an owner state without updating memory of the system.

25. (original) The system of claim 24, further comprising:

means for broadcasting from the first node a migratory read (XREADM) request for data;

means for providing the data from the second node to the first node in response to the XREADM request, the modified state being associated with the data at the second node, the shared state being associated with the data at the first node in response to the first node receiving the data from the second node; and

means for transitioning the modified state associated with the data at the second node to an invalid state without updating memory of the system.

26. (previously presented) The system of claim 25, further comprising means for selecting one of the XREADM request and XREADN request to broadcast from the first node.

27. (previously presented) The system of claim 25, further comprising means for predictively selecting one of the XREADM request and XREADN request to broadcast from the first node.

28. (original) A method comprising:

broadcasting a non-migratory request for data from a first node to other nodes of an associated system;

providing a shared copy of the data from a second node to the first node in response to the non-migratory request;

transitioning a state associated with the data at the second node from a modified state to an owner state in response to the non-migratory request; and

transitioning a state associated with the data at the first node to a shared state in response to receiving the shared copy of the data from the second node.

29. (original) The method of claim 28, further comprising:

broadcasting a migratory request for the data from the first node to other nodes of the associated system;

providing an ownership data response from the second node to the first node in response to the migratory request;

transitioning the state associated with the data at the second node from a modified state to an invalid state data in response to the migratory request; and

transitioning the state associated with the data at the first node to a dirty state in response to receiving the ownership data response from the second node.

30. (original) The method of claim 29, further comprising selecting one of the migratory request and the non-migratory request to broadcast from the first node.

31. (original) The method of claim 29, further comprising predictively selecting one of the migratory request and the non-migratory request to broadcast from the first node.

32. (original) A computer system comprising a cache coherency protocol that is operative to permit migration of data to a cache associated with a source processor from a cache associated with a target processor when a migratory request is issued from the source processor, the protocol being further operative to prevent migration of the data to the cache associated with the source processor from the cache associated with the target processor when a non-migratory request is issued from the source processor.

IX. Evidence Appendix

None

X. Related Proceedings Appendix

None